arXiv:cond-mat/0312317v1 [cond-mat.stat-mech] 12 Dec 2003

e-print published at http://www.arxiv.org/abs/cond-mat/0312317 December 2003

Estimating Driving Forces of Nonstationary Time Series with Slow Feature Analysis

Laurenz Wiskott

Institute for Theoretical Biology, Humboldt-University Berlin
Invalidenstraße 43, D-10115 Berlin, Germany
http://itb.biologie.hu-berlin.de/~wiskott
l.wiskott@biologie.hu-berlin.de

Abstract

Slow feature analysis (SFA) is a new technique for extracting slowly varying features from a quickly varying signal. It is shown here that SFA can be applied to nonstationary time series to estimate a single underlying driving force with high accuracy up to a constant offset and a factor. Examples with a tent map and a logistic map illustrate the performance.

1 Introduction

Nonlinear time series analysis generally assumes stationarity (see Casdagli, 1997; Schreiber, 1999, for an overview). However, many time series are actually nonstationary for various reasons, such as temperature drift in the experimental setup, decreasing reservoirs in (bio)chemical reactors or ecological systems, global warming in climate data, varying heart rate in cardiology, or vibrato/tremor in vocal fold vibrations. Such nonstationarities can be modeled by underlying parameters, referred to as driving forces, that change the dynamics of the system smoothly on a slow time scale or abruptly but rarely, e.g. if the dynamics switches between different discrete states.

If a test reveals that a time series is nonstationary (Witt et al., 1998), one can still apply methods of stationary time series analysis if one determines sections where the driving force has similar values and analyzes these sections as one stationary time series. This can be done by first slicing the time series into windows of equal size and then grouping the windows based on similarity measures of the dynamics (Manuca & Savit, 1996). This division of the time series can be avoided by the technique of overembedding. If the embedding dimension is sufficiently high, then similar embedding vectors automatically belong to similar values of the driving forces and all available data can be used for analysis, such as forecasting or nonlinear noise reduction (Hegger et al., 2000).

However, in some cases, e.g. in the analysis of EEG data, one is particularly interested in revealing the driving forces themselves, which is in principle only possible up to an invertible transformation. One standard method for visualizing driving forces is the recurrence plot (Eckmann et al., 1987), but it is often difficult to interpret. Methods have been developed to estimate driving forces based on the finding that the recurrence plot of a time series is similar to the recurrence plot of its underlying driving forces (Casdagli, 1997; Schreiber, 1999). Another technique for the reconstruction of driving forces has been presented in (Verdes et al., 2001). Here I present an alternative approach based on slow feature analysis, a new technique developed in the field of theoretical neurobiology.

2 Slow Feature Analysis

Slow feature analysis (SFA) has been originally developed in context of an abstract model of unsupervised learning of invariances in the visual system of vertebrates (Wiskott, 1998) and is described in detail in (Wiskott & Sejnowski, 2002).

The general objective of SFA is to extract slowly varying features from a quickly varying signal. For a scalar output signal it can be formalized as follows. Let $\boldsymbol{x} = \boldsymbol{x}(t)$ be an N-dimensional input signal where t indicates time and $\boldsymbol{x} = [x_1, ..., x_N]^T$ is a vector. Find the input-output function $g(\boldsymbol{x})$ that generates a scalar output signal

$$y(t) := g(\boldsymbol{x}(t)) \tag{1}$$

with most slowly temporal variation possible as measured by the variance of the time derivative:

$$\text{minimize} \quad \langle \dot{y}^2 \rangle \tag{2}$$

with $\langle \cdot \rangle$ indicating the temporal mean. For convenience and to avoid the trivial constant solution the output signal has to meet the following constraints:

$$\langle y \rangle = 0 \quad \text{(zero mean)}\,, \tag{3}$$
$$\langle y^2 \rangle = 1 \quad \text{(unit variance)}\,. \tag{4}$$

This is an optimization problem of variational calculus and as such difficult to solve. However, if we constrain the input-output function to be a linear combination of some fixed and possibly nonlinear basis functions, the problem becomes tractable and can be solved in the following way. Let $\boldsymbol{h}' = \boldsymbol{h}'(\boldsymbol{x})$ be a vector of some fixed basis functions. To be concrete assume $\boldsymbol{h}'$ contains all monomials of degree 1 and 2. Applying $\boldsymbol{h}'$ to the input signal $\boldsymbol{x}(t)$ yields the nonlinearly expanded signal $\boldsymbol{z}'(t)$:

$$\boldsymbol{h}'(\boldsymbol{x}) := [x_1,\, x_2,\, ...,\, x_N,\, x_1^2,\, x_1x_2,\, ...,\, x_N^2]^T, \tag{5}$$
$$\boldsymbol{z}'(t) := \boldsymbol{h}'(\boldsymbol{x}(t))\,. \tag{6}$$

Assume $g(\boldsymbol{x})$ is a linear combination of the basis functions plus a constant c, generating the output signal $y(t) = g(\boldsymbol{x}(t))$:

$$g(\boldsymbol{x}) := c + \boldsymbol{a}'^T \boldsymbol{h}'(\boldsymbol{x})\,, \tag{7}$$
$$y(t) = c + \boldsymbol{a}'^T \boldsymbol{z}'(t)\,. \tag{8}$$

Since $\boldsymbol{h}'$ contains all monomials of degree 1 and 2, $g(\boldsymbol{x})$ can be any polynomial of degree 2, if the constant c and the coefficient vector $\boldsymbol{a}'$ are chosen correspondingly.

For reasons that become clear below, it is convenient to sphere (or whiten) the expanded signal and transform the basis functions accordingly:

$$\boldsymbol{h}(\boldsymbol{x}) := \boldsymbol{S}\,(\boldsymbol{h}'(\boldsymbol{x}) - \langle \boldsymbol{z}' \rangle)\,, \tag{9}$$
$$\boldsymbol{z}(t) := \boldsymbol{S}\,(\boldsymbol{z}'(t) - \langle \boldsymbol{z}' \rangle)\,, \tag{10}$$

with the sphering matrix $\boldsymbol{S}$ chosen such that the signal components have a unit covariance matrix, i.e. $\langle \boldsymbol{z}\boldsymbol{z}^T \rangle = \boldsymbol{I}$, which can be easily done with the help of principal component analysis (PCA). The signal components have also zero mean, $\langle \boldsymbol{z} \rangle = \boldsymbol{0}$, since the mean values have been subtracted.

The input-output function and the output signal can now be written in these transformed basis functions and sphered expanded signal, respectively:

$$g(\boldsymbol{x}) = \boldsymbol{a}^T \boldsymbol{h}(\boldsymbol{x})\,, \tag{11}$$
$$y(t) = \boldsymbol{a}^T \boldsymbol{z}(t)\,. \tag{12}$$

This is a form in which the constraints (3, 4) are particularly easy to meet, since we find for any coefficient vector $\boldsymbol{a}$ with norm 1,

$$\langle y \rangle = \boldsymbol{a}^T \underbrace{\langle \boldsymbol{z} \rangle}_{=\boldsymbol{0}} = 0\,, \tag{13}$$

$$\langle y^2 \rangle = \boldsymbol{a}^T \underbrace{\langle \boldsymbol{z}\boldsymbol{z}^T \rangle}_{=\boldsymbol{I}} \boldsymbol{a} = \boldsymbol{a}^T \boldsymbol{a} = 1\,, \tag{14}$$

which also motivates why constant c has been dropped in (11, 12).

Thus the optimization problem reduces to finding the normalized coefficient vector $\boldsymbol{a}$ that minimizes

$$\langle \dot{y}^2 \rangle = \boldsymbol{a}^T \langle \dot{\boldsymbol{z}} \dot{\boldsymbol{z}}^T \rangle \boldsymbol{a} \,, \tag{15}$$

which is obviously the normalized eigenvector of the time-derivative covariance matrix $\langle \dot{\boldsymbol{z}} \dot{\boldsymbol{z}}^T \rangle$ with the smallest eigenvalue, which again can be easily found by PCA. The output signal $y(t)$ is then given by (12). Notice that $y(t)$ is usually uniquely determined up to the sign.

In summary SFA consists of basically four steps: (i) expand the input signal with some set of fixed possibly nonlinear functions; (ii) sphere the expanded signal to obtain components with zero mean and unit covariance matrix; (iii) compute the time derivative of the sphered expanded signal and determine the normalized eigenvector of its covariance matrix with the smallest eigenvalue; (iv) project the sphered expanded signal onto this eigenvector to obtain the output signal. The input-output function is given by (11) but is not needed here, since we are only interested in the extracted output signal.

In practice one has to work with time series $\boldsymbol{x}_i$ instead of continuous signals $\boldsymbol{x}(t)$, but the transfer of the algorithm described above to time series is straight forward. The time derivative is simply computed as the difference between successive data points without wrap-around assuming a constant sampling spacing Δt.

The description of SFA given above assumes that all monomials of degree 1 and 2 are used for signal expansion. However, any other set of basis functions could be used as well, e.g. monomials of higher degree or radial basis functions. It is also straight forward to extend SFA to more than one output signal component, in which case additional components should be uncorrelated to earlier ones. Under this constraint additional components can be computed by using the normalized eigenvectors of (15) with the next larger eigenvalues (Wiskott & Sejnowski, 2002). SFA requires that all components vary to some extent. Thus if there are constant components in the input signal or the expanded signal, these are simply discarded.

3 Examples

In the following I will present two examples with time series w_i derived from a tent map and a logistic map to illustrate the properties of SFA. I have also done simulations with the Lorenz System, but results were very unreliable and sensitive to parameter variations and noise. The underlying driving force will always be denoted by γ and may vary between -1 and 1 either smoothly or rarely, but with a comparable time scale of variation (as defined by the variance of its time derivative (2)).

Taking the time series w_i directly as an input signal would not give SFA enough information to estimate the driving force, because SFA considers only data (and its derivative) from one time point at a time. Thus it is necessary to generate an embedding-vector time series as an input. Here embedding vectors at time point i with delay τ and dimension m are defined by

$$\boldsymbol{x}_i \;:=\; \left[w_{i-\tau(m-1)/2},\, w_{i-\tau((m-1)/2-1)},\, ...,\, w_{i+\tau(m-1)/2}\right]^T , \tag{16}$$

for scalar w_i and odd m. The definition can be easily extended to even m, which requires an extra shift of the indices by $\tau/2$ or its next lower integer to center the used data points at i. Centering the embedding vectors results in an optimal temporal alignment between estimated and true driving force.

The following simulations are based on 6000 data points each and were done with Matlab (Release 13).

3.1 Tent Map

As a first example consider a time series generated by an iterated tent map (Casdagli, 1997)

$$\hat{\gamma}_i \;:=\; (\gamma_i + 1)/2 \tag{17}$$

$$w_{i+1} \;=\; \begin{cases} +2w_i + \hat{\gamma}_i & \text{if} \quad 0 \le w_i \le \frac{1}{2}(1-\hat{\gamma}_i) \\ -2w_i - \hat{\gamma}_i + 2 & \text{if} \quad \frac{1}{2}(1-\hat{\gamma}_i) < w_i \le \frac{1}{2}(2-\hat{\gamma}_i) \\ +2w_i + \hat{\gamma}_i - 2 & \text{if} \quad \frac{1}{2}(2-\hat{\gamma}_i) < w_i \le 1 \end{cases} , \tag{18}$$

which maps the interval $[0,1]$ onto itself with a functional form "$\bigwedge$" for $\gamma = -1$ (or $\hat{\gamma} = 0$). The parameter γ shifts this wedge cyclically to the left until it becomes a "$\bigvee$" for $\gamma = \hat{\gamma} = +1$. Figure 1 shows the true driving force, the time series, and the estimated driving force with $m = 10$, $\tau = 1$, and third order polynomials for SFA. The correlation between true and estimated driving force is $r = 0.96$. Note that the scale and offset of the estimated driving force are arbitrarily fixed by the constraints and that the sign is random. The axes were therefore chosen such that the curves in the bottom graphs of this and the following figures are optimally aligned with each other.

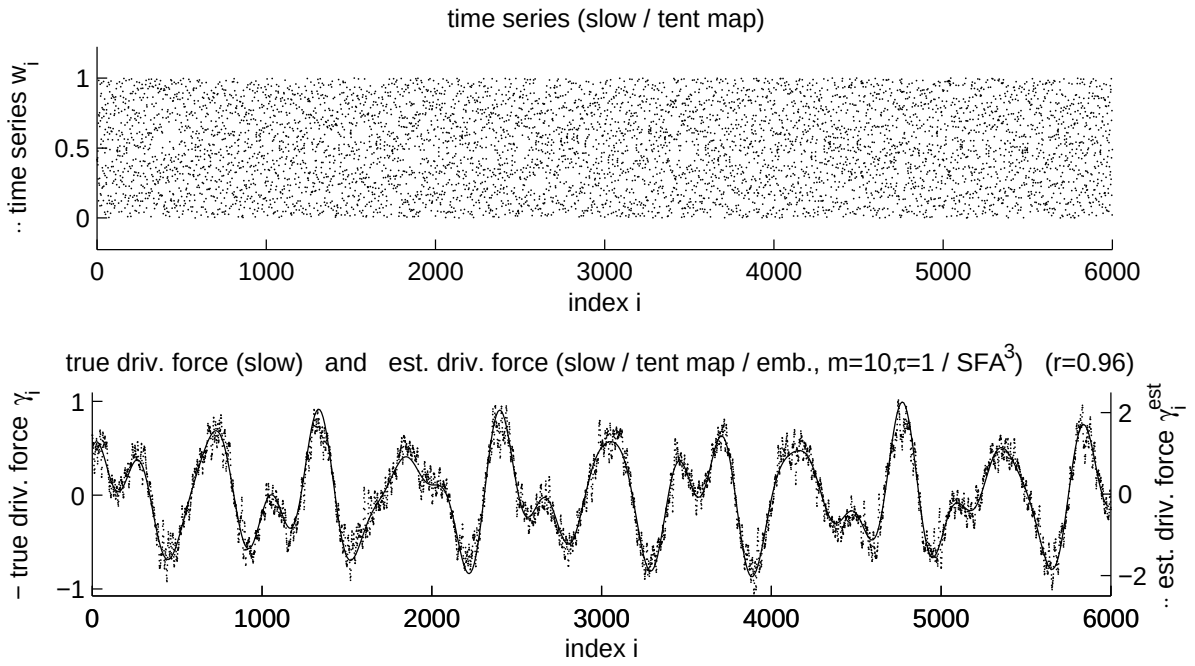

Figure 1: True slowly varying driving force γ_i (bottom, solid line), time series w_i (top) derived from the driving force with the tent map, and estimated driving force γ_i^{est} (bottom, dots). The correlation between true and estimated driving force is $r = 0.96$.

3.2 Logistic Map

As a second example consider a time series derived from a logistic map

$$w_{i+1} = (3.6 + 0.4\gamma_i) w_i (1 - w_i)\,, \tag{19}$$

which maps the interval $[0, 1]$ onto the interval $[0, 0.9 + 0.1\gamma_i]$ and has the shape of an upside-down parabola crossing the abscissa at 0 and 1. Parameter γ governs the height of the parabola. Figure 2 shows the true driving force, the time series, and the estimated driving force with $m = 3$, $\tau = 1$, and second order polynomials for SFA. The correlation between true and estimated driving force is $r = 0.997$.

4 Technical Remarks

Choice of Parameters SFA is basically parameter-free except for the general choice of the class of nonlinear basis functions. The only other parameters to choose in this method are the dimension m and the time delay τ of the embedding vectors. The value of $\langle \dot{y}^2 \rangle$ is a fairly reliable indicator for a good choice of values for m and τ. The smaller $\langle \dot{y}^2 \rangle$ the better, since there is no trivial way of achieving small $\langle \dot{y}^2 \rangle$-values (except if the number of basis functions comes close to the number of data points). I typically test a certain range of τ-values and successively increase the m-value until performance (measured in terms of $\langle \dot{y}^2 \rangle$ or r) is satisfactory.

Rarely Varying Driving Forces If SFA is designed to extract slowly varying features from a signal, how about driving forces that apparently violate this assumption heavily, e.g. if they vary abruptly and jump between different discrete values? The definition of slowness given in (2) does actually not depend on the smoothness of the output signal, it may equally well be a signal that switches abruptly between different discrete values as long as the jumps occur rarely enough to lead to the same variance of the time derivative. Whether the order of the values can be recovered depends on the dynamics of the driving force. SFA will tend to order the values such that jumps occur between nearby values of the estimated driving force. Thus, if the original values are -0.5, 0, and +0.5 and there are many direct jumps between -0.5 and +0.5, SFA might map the three values onto -1, +1, and 0, respectively, thereby changing the order of the second and

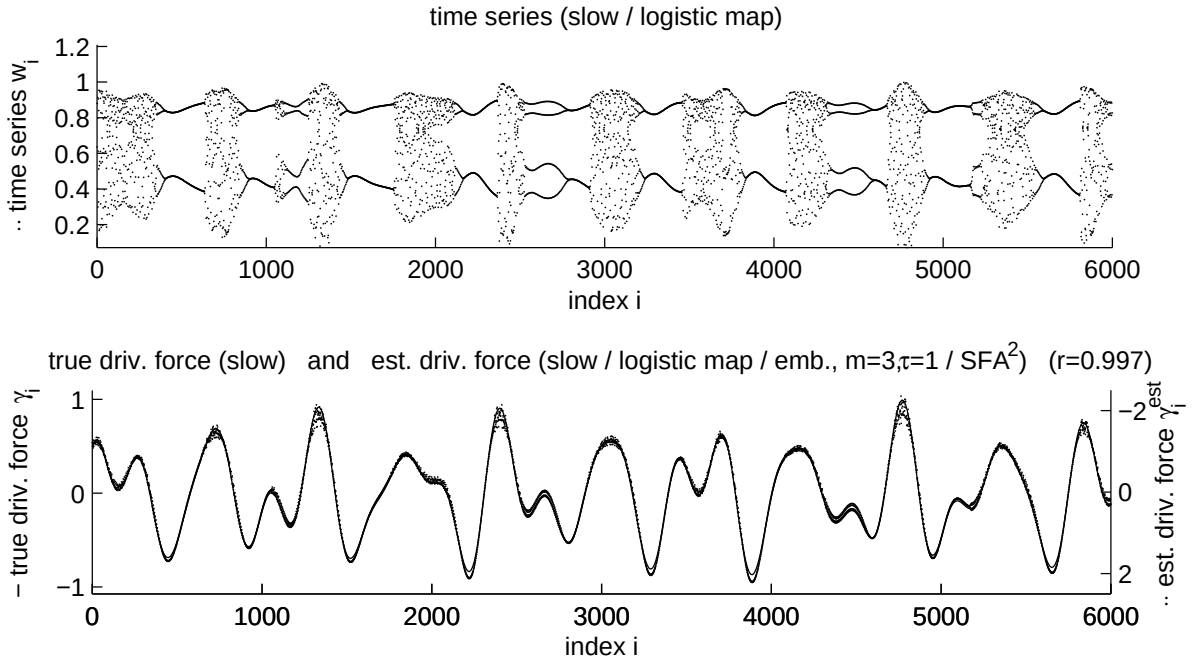

Figure 2: True slowly varying driving force γ_i (bottom, solid line), time series w_i (top) derived from the driving force with the logistic map, and estimated driving force γ_i^{est} (bottom, dots). The correlation between true and estimated driving force is $r = 0.997$.

third value. Figures 3 and 4 show results for the two systems if a rarely instead of a slowly and smoothly varying driving force is used.

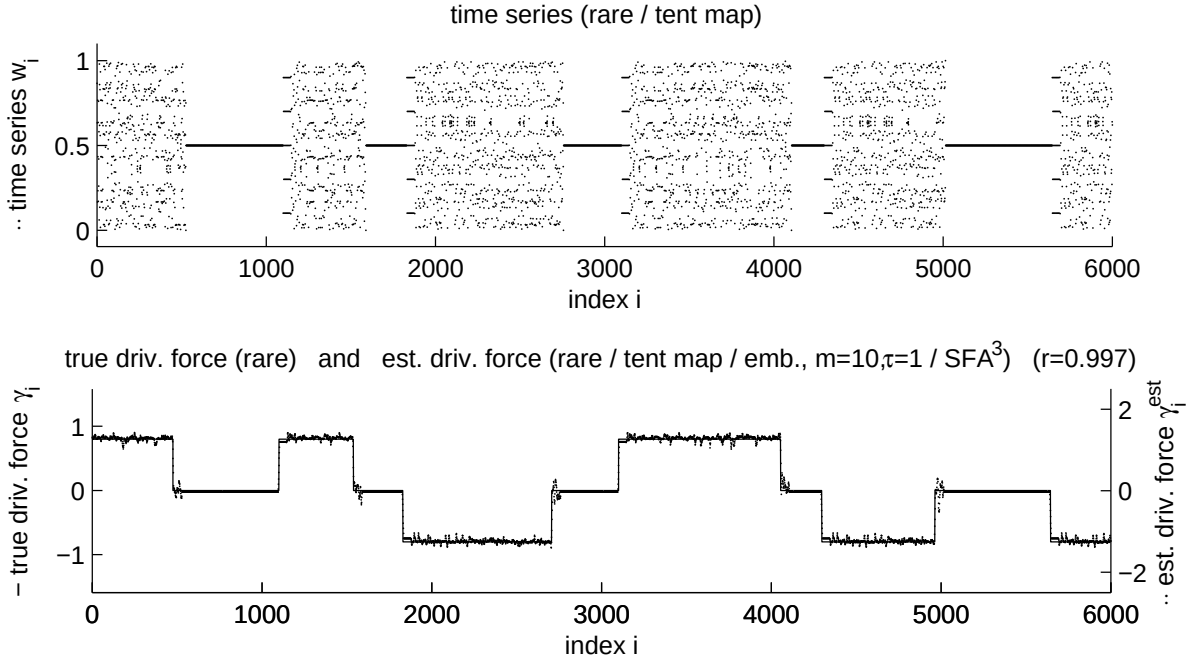

Figure 3: True rarely varying driving force γ_i (bottom, solid line), time series w_i (top) derived from the driving force with the tent map, and estimated driving force γ_i^{est} (bottom, dots). The correlation between true and estimated driving force is $r = 0.997$.

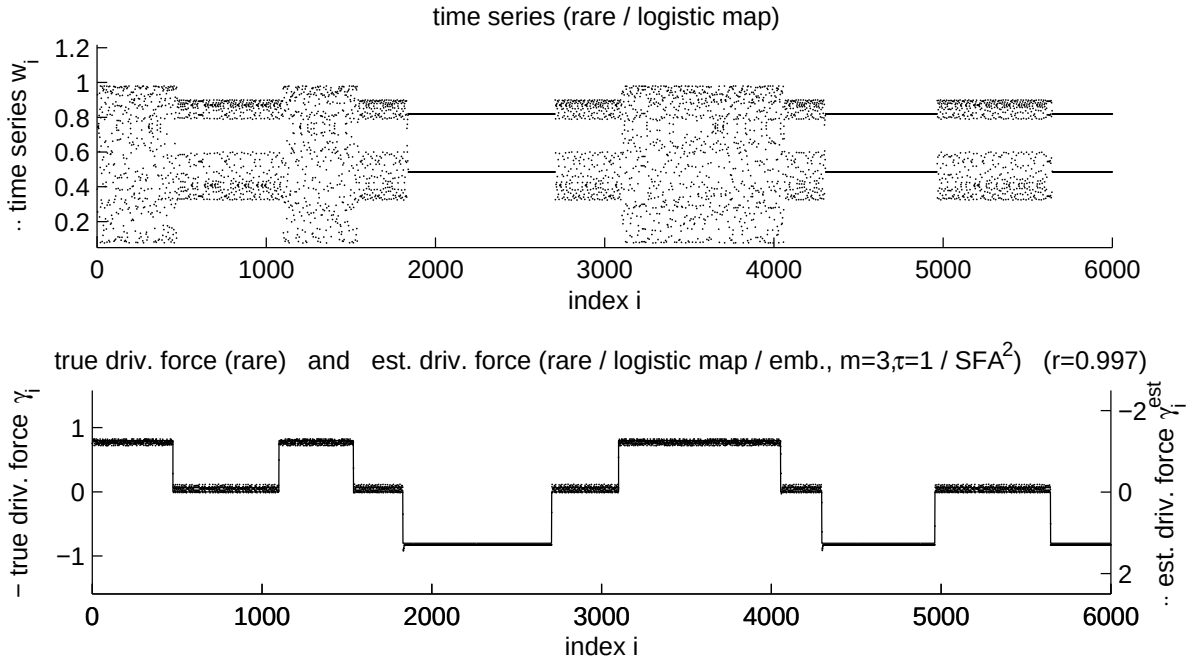

Figure 4: True rarely varying driving force γ_i (bottom, solid line), time series w_i (top) derived from the driving force with the logistic map, and estimated driving force γ_i^{est} (bottom, dots). The correlation between true and estimated driving force is $r = 0.997$.

High-Dimensional Input Data Probably the most severe limitation of SFA is the fact that the number of monomials (or any other basis functions) grows quickly with the dimensionality of the embedding vectors (curse of dimensionality). One therefore might run out of computer memory before a high enough m-value is reached. However, higher-dimensional problems could be dealt with in a hierarchical fashion by breaking the embedding vectors into smaller parts which are first analyzed separately and the results of which are then combined for a final analysis. Applying this hierarchical scheme to 65-dimensional input vectors has been demonstrated in (Wiskott & Sejnowski, 2002).

Accuracy of the Estimated Driving Force It might be surprising that in the examples above SFA is able to estimate the driving forces with such an accuracy up to a factor and a constant offset, even though the estimation is undetermined up to any invertible transformation, not only scaling and shift. One reason for that is that the driving forces used are relatively slow already and could probably not be improved much by an invertible nonlinear transformation. However, even if that were not the case one might hope that in practice for the lower-dimensional embedding vectors ($m = 10$ instead of, e.g., $m = 50$) SFA has to find a relatively simple input-output function $g(\boldsymbol{x})$, which is more likely to preserve the exact shape of the driving-force curve than a more complicated one. Further experiments are needed to verify this.

Noise Sensitivity How sensitive is the method to noise? One might suspect that SFA is quite sensitive to noise, since eigenvectors of smallest eigenvalues are used. However, focusing on small eigenvalues is only done after sphering and taking the time derivative, so that small but quickly varying noise components typically have large eigenvalues and are discarded by SFA. Graceful degradation with noise was also found in (Wiskott & Sejnowski, 2002). Thus, using the eigenvectors with smallest eigenvalues in itself does not induce noise sensitivity. I have done simulations with the examples presented here by adding Gaussian white noise to the signals before embedding. Adding 10%, 20%, and 50% noise to the tent map time series reduced the correlation between true and estimated driving force from $r = 0.96$ to about 0.94, 0.90, and 0.71, respectively; adding 2%, 5%, and 10% noise to the logistic map time series reduced the correlation from $r = 0.997$ to about 0.97, 0.87, and 0.70, respectively. Thus we see, that in these examples the method is in fact fairly robust with respect to noise.

Multiple Driving Forces SFA can be easily extended to the extraction of multidimensional output signals (Wiskott & Sejnowski, 2002), which could be used to estimate multiple driving forces. However, this might require longer time series and higher-dimensional embedding vectors. If the different driving forces are not clearly separated by different time scales, SFA might only be able to estimate them up to a linear mixing transformation. In that case independent component analysis (ICA) might be able to separate them, if they are statistically independent and not more than one is Gaussian (see Hyvärinen, 1999, for an overview of ICA methods).

5 Conclusion

In this paper I have demonstrated that slow feature analysis (SFA) can be applied to the problem of estimating a driving force of a nonstationary time series. The estimates are fairly accurate except for a scaling factor and a constant offset, which cannot be extracted in general. The method works for slowly as well as rarely varying driving forces, although in the latter case the order of the different discrete values might not be recoverable. The slowness principle also provides robustness to noise, which is typically quickly varying and therefore suppressed by SFA as far as possible.

There are still a number of open questions. The next steps of investigation will have to include a comparison with standard methods, such as the technique of recurrence plots, and an exploration of the conditions under which SFA can be applied with similar success as demonstrated here. It is also necessary to apply SFA to real world data. This has been successfully done in the context of learning receptive field properties of the visual cortex based on natural image sequences (Berkes & Wiskott, 2003), but that was not for extracting driving forces. In any case, since SFA works very differently from other techniques that have been applied to the estimation of driving forces, one can hope that it at least complements these other techniques in some cases.

6 Acknowledgment

I am grateful to Hanspeter Herzel and Isao Tokuda for useful hints and Hanspeter Herzel also for critically reading the manuscript. This work has been supported by the VolkswagenStiftung.

References

Berkes, P. and Wiskott, L. (2003). Slow feature analysis yields a rich repertoire of complex-cell properties. Cognitive Sciences EPrint Archive (CogPrints) 2804, http://cogprints.ecs.soton.ac.uk/archive/00002804/.

Casdagli, M. C. (1997). Recurrence plots revisited. *Physica D: Nonlinear Phenomena*, 108(1-2):12–44.

Eckmann, J. P., Oliffson Kamhorst, S., and Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9):973–977.

Hegger, R., Kantz, H., Matassini, L., and Schreiber, T. (2000). Coping with nonstationarity by overembedding. *Physical Review Letters*, 84(18):4092–4095.

Hyvärinen, A. (1999). Survey on independent component analysis. *Neural Computing Surveys*, 2:94–128.

Manuca, R. and Savit, R. (1996). Stationarity and nonstationarity in time series analysis. *Physica D: Nonlinear Phenomena*, 99(2-3):134–161.

Schreiber, T. (1999). Interdisciplinary application of nonlinear time series methods. *Physics Reports*, 308(1):1–64.

Verdes, P. F., Granitto, P. M., Navone, H. D., and Ceccatto, H. A. (2001). Nonstationary time-series analysis: Accurate reconstruction of driving forces. *Physical Review Letters*, 87(12):124101.

Wiskott, L. (1998). Learning invariance manifolds. In *Proc. 5th Joint Symp. on Neural Computation, San Diego*, pages 196–203. University of California, San Diego.

Wiskott, L. and Sejnowski, T. (2002). Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770.

Witt, A., Kurths, J., and Pikovsky, A. (1998). Testing stationarity in time series. *Physical Review E*, 58(2):1800–1810.